

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JUL 2012		2. REPORT TYPE		3. DATES COVERED 00-00-2012 to 00-00-2012	
4. TITLE AND SUBTITLE Integrating Hard and Soft Information Sources for D2D Using Controlled Natural Language				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) US Army Research Laboratory, Adelphi, MD, 20783-1138				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Presented at the 15th International Conference on Information Fusion held in Singapore on 9-12 July 2012. Sponsored in part by Office of Naval Research and Office of Naval Research Global.					
14. ABSTRACT We introduce an approach to integrating access to hard and soft information sources to provide better exploitation of all available sources in the context of coalition data-to-decision (D2D) chains. In terms of hard (sensor-based) sources we show how intelligence, surveillance, and reconnaissance (ISR) assets can be represented at a relatively high level in controlled natural language, and how this allows the automatic assignment of sensing assets to D2D tasks. We demonstrate how the use of Controlled English (CE) ??? a type of controlled natural language designed to be readable by a native English speaker whilst representing information in a structured, unambiguous form ??? supports the informed sharing of D2D tasks and assets between collaborating users in a coalition environment. Moreover, we show how CE can be used in the automatic extraction of information from unstructured and semi-structured text information sources, providing us with a uniform way to integrate these soft sources with the aforementioned hard sources.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Integrating Hard and Soft Information Sources for D2D Using Controlled Natural Language

Alun Preece*, Diego Pizzocaro*, Dave Braines[†], David Mott[†], Geeth de Mel^{‡§} and Tien Pham[§]

*School of Computer Science and Informatics, Cardiff University, Cardiff, UK; contact email: A.D.Preece@cs.cf.ac.uk

[†]Emerging Technology Services, IBM United Kingdom Ltd, Hursley Park, Winchester, UK

[‡]IBM TJ Watson Research Center, Hawthorne NY, USA

[§]US Army Research Laboratory, Adelphi MD, USA

Abstract—We introduce an approach to integrating access to hard and soft information sources to provide better exploitation of all available sources in the context of coalition data-to-decision (D2D) chains. In terms of hard (sensor-based) sources we show how intelligence, surveillance, and reconnaissance (ISR) assets can be represented at a relatively high level in controlled natural language, and how this allows the automatic assignment of sensing assets to D2D tasks. We demonstrate how the use of Controlled English (CE) — a type of controlled natural language designed to be readable by a native English speaker whilst representing information in a structured, unambiguous form — supports the informed sharing of D2D tasks and assets between collaborating users in a coalition environment. Moreover, we show how CE can be used in the automatic extraction of information from unstructured and semi-structured text information sources, providing us with a uniform way to integrate these soft sources with the aforementioned hard sources.

I. INTRODUCTION

The term *data-to-decision* (D2D) characterises decision-making domains where data and sources of data are plentiful, but it is difficult to assemble rapidly the right set of data to facilitate making decisions.¹ D2D emphasises the collection and fusion of actionable information, to provide a clear picture of options, threats, and consequences [1], [2]. In this context, decision makers may be at any level in an organisation, from high level leaders at the centre of an information network, to low level operatives on the edge of the network. In particular, D2D aims to empower individuals on the edge who, prior to the widespread provision of mobile information and communication platforms, have not traditionally been able to benefit from the best-available actionable information. In domains such as emergency response, policing, and military operations, empowering such individuals is important since they are the ones whose actions will have a direct effect on the unfolding situation. Because the situation evolves rapidly, the information-provisioning infrastructure that supports D2D activity must be agile, being responsive to changes in the decision-maker's needs and the availability of relevant sources.

We observe that an agile D2D information infrastructure requires information to flow in two directions:

- Forwards from data to decision: a decision-maker needs to take decisions based on actionable information, pro-

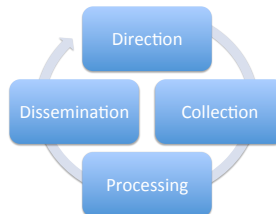


Fig. 1. The direction-collection-processing-dissemination (DCPD) cycle

cessed from data collected by sensors (for example, imagery or audio data) or retrieved from other sources (for example, eyewitness reports, newsfeeds).

- Backwards from the intended decision to relevant information assets: a decision-maker needs to determine what kinds of information will help them achieve their overall intent, and thereby identify suitable assets.

It has commonly been observed (for example, in [3]) that these flows form part of a continuous cycle. The stages in this cycle correspond to stages in the intelligence process, which in the UK are referred to as direction, collection, processing, and dissemination (DCPD). Figure 1 shows the DCPD cycle. The stages in this cycle are defined as follows²:

- *Direction*: the determination of information needs and sources relevant to an intended decision.
- *Collection*: the acquisition of data from relevant sources (for example, sensors, databases, or feeds).
- *Processing*: turning collected data into usable information (for example, by addition of context, indexing, fusion, or filtering).
- *Dissemination*: the transmission of processed information to decision makers.

The process is considered a cycle because received information typically provokes further information needs; for example, to obtain more detailed observations on a newly-detected object of interest, or to explore the feasibility of possible responses to a perceived threat.

²A US variant of the DCPD cycle (called TCPED) refers to direction as *tasking* and breaks the DCPD processing step into two stages, *processing* and *exploitation*, where the former is essentially “pre-processing” to put data into a usable form, and the latter involves putting the information into the context of a particular decision.

¹<http://www.acq.osd.mil/chieftechnologist/areas/dtd.html>

The D2D context demands that the DCPD process be implemented in a highly agile manner, which implies that it should be automated to the greatest extent possible. Software can assist in many ways, from the identification of relevant sources, to the automatic generation of queries and sensor tasking requests, to the composition and invocation of useful information-processing services, to the selection of appropriate dissemination mechanisms which take into account the capabilities of an end-user's (mobile) device. Infrastructure to support D2D activities must be resilient in the face of rapidly-changing information needs and availability of assets (for example, sensors may fail or sources may go offline). Many of the technical elements required for an agile and resilient D2D-supporting system are discussed in [4]. Any system to support D2D needs to be capable of exploiting both hard and soft sources. *Hard* sources are those derived from physics-based sensing (for example, data collected from video, acoustic, or seismic sensors) while *soft* sources originate from humans (for example, eyewitness reports, newsfeeds, or text messages).

In this paper, we propose the use of a *controlled natural language* (CNL) as an enabling component of an agile infrastructure to support D2D activities. A CNL is a subset of a natural language, commonly English, with restricted syntax and vocabulary. Often they are used to provide an information representation that is easily machine processable (with low complexity and no ambiguity) while also being human-readable (see, for example, [5]). We show how a form of Controlled English (CE) can be used to represent elements necessary for the automation of the DCPD cycle, as follows:

- A. *Information needs*: a machine-processable representation of a decision-maker's information needs is a prerequisite for software assistance in determining what collection assets are relevant, as part of the direction stage.
- B. *Asset capabilities*: metadata describing what information an asset can provide is a prerequisite for asset selection as part of the direction stage; this can be seen as a kind of "matchmaking" process against the information needs (A).
- C. *Information products*: in order for data collected from assets (either hard or soft) to be processed further (including fusion) it needs to be transformed into a machine-processable form, consistent with the metadata specified for the asset (B); this supports the processing stage, and the ultimate delivery of data to meet the original information needs (A) in the dissemination stage.

The rest of the paper is structured as follows: Section II introduces a detailed but reasonably generic surveillance vignette that we will use to illustrate the CNL-based approach. Section III summarises a number of system components from previous work that help in automating the DCPD loop to support agile D2D activities. Sections IV and V describe how CNL can assist in the identification and exploitation of hard and soft sources, respectively, and the fusion of information from a combination of hard and soft sources. Section VI provides a detailed walkthrough of the vignette with CNL examples. Section VII summarises and concludes the paper.

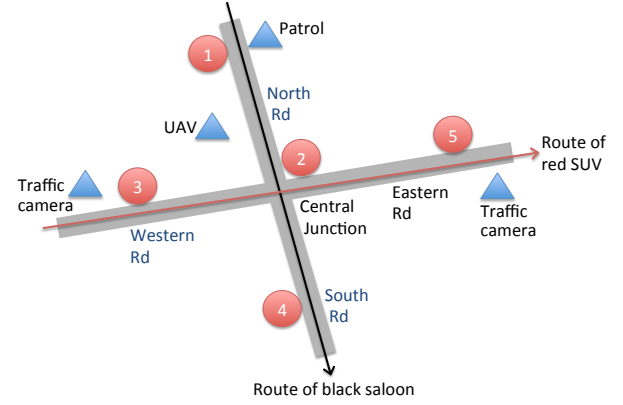


Fig. 2. A surveillance D2D vignette with hard and soft information fusion

II. AN ILLUSTRATIVE VIGNETTE

We introduce a small vignette which will be used in later sections to illustrate the use of our CNL-based approach to supporting D2D activities. An area of interest around the intersection of four roads is shown in Figure 2. The locations of various assets — sources of both hard and soft information — are marked by triangles. The passage of two vehicles causes a sequence of events, shown as numbered points on the map, to unfold as follows:

- 1) A patrol on North Road reports a suspicious black saloon car, vehicle registration ABC123, moving south. A database query reveals that this vehicle is known to be associated with a high value target, John Smith. A request is issued to track the location of the vehicle. An unmanned aerial vehicle (UAV) is assigned to this task.
- 2) The UAV locates and tracks the black saloon as it heads south on North Road. The UAV reports that the vehicle stops near Central Junction. An analyst is alerted of this, and requests imagery from the UAV. They find that the black saloon has stopped by the roadside next to a red SUV. The analyst indicates that the red SUV is an object of interest. The two vehicles now depart the junction, the saloon heading south onto South Road, and the SUV heading east on Eastern Road.
- 3) The analyst's indication that the red SUV is now of interest causes a recent report to be retrieved from a camera system on Western Road: a red SUV passed the camera recently; license plate recognition software determined that its registration is XYZ789. This identification is now associated with the SUV from Central Junction with a high degree of certainty, given the recency of the report and the fact that no other similar SUVs passed by.
- 4) As the saloon and SUV head south and east respectively, decisions need to be taken on whether and how to track their movements. The only available assets in the area are the UAV and a traffic camera system on Eastern Road. Both are capable of locating the vehicles, though the camera system can only do so in a limited area. In the event, the UAV is tasked to continue following the saloon

in the South Road area.

- 5) The camera system on Eastern Road is tasked to issue an alert on identifying a red SUV with license plate XYZ789. As this only covers part of the road, local law enforcement in the Eastern Road area are also alerted to look out for this vehicle.

In the next section, we summarise elements of the D2D infrastructure required to provide automated support for the steps in this vignette.

III. D2D INFRASTRUCTURE AND SERVICES

The following is not intended to be an exhaustive list of required infrastructure and services, but rather to identify a number of key elements and assumptions. Following [6], [7], we assume the existence of a service-oriented architecture:

- There is an online catalogue of available assets, described in terms of their capabilities, and a knowledge base that determines the selection of assets suitable for a particular task (as part of the direction DCPD stage). In the vignette for example, we need to know that the UAV is capable of determining the location of a vehicle with particular features (type and colour). Similarly, the cameras on Western and Eastern Roads can identify vehicles by license plate. Visibility of the asset catalogue — and access to the assets as described below — is governed by access policies, as highlighted in [4] but not described further here. Note that both hard and soft assets can be catalogued in terms of capabilities (the kinds of information they can provide). We assume there is a procedure for allocating assets to tasks where there are competing demands [8].
- Assets are “wrapped” as network services so they can deliver data in terms of typed information feeds. In our example, the UAV is tasked to deliver a series of location reports on the target vehicle, while the cameras are tasked to deliver reports on vehicle features (including type, colour, and registration if visible). The UAV itself will generate “raw” data and wrapper services will process this into typed information (location reports). Some assets may be capable of delivering multiple feeds; for example, both the UAV and cameras also allow the original “raw” imagery data corresponding to a processed report to be retrieved (subject to available network bandwidth). Note that these requirements address aspects of the collection, processing, and dissemination DCPD stages.
- Soft sources are “wrapped” in similar ways to assets. Here, the original data will be unstructured (or at best semistructured) text, and a key challenge is to process this to yield usable information in an agile manner. Depending on the complexity of the text, it may be possible to automatically extract some elements, either online (in real-time as the messages come in) or offline (typically by collecting them in a database and processing them, often with reference to other sources to aid analysis). Again, these features cover aspects of the collection, processing, and dissemination DCPD stages.

- There is a set of ontologies describing the entities and relationships in the domain of interest, against which we can capture instance data. These ontologies cover not only objects of interest in the world (for example, vehicles, people, places, times) but also elements of the intelligence cycle itself (for example, the various kinds of assets and tasks). Further details on ontologies associated with the intelligence, surveillance, and reconnaissance domain are given in [6]. The set of ontologies must be extensible, modular, and capable of being interlinked in flexible ways. They support all stages of the DCPD cycle, as we will show in the next two sections.

IV. USING CNL IN EXPLOITATION OF HARD SOURCES

As noted above, our approach depends on the existence of a set of ontologies, representing elements of the domain of interest as well as the DCPD process itself. In the CNL-based approach, we use Controlled English (CE) to express these. The purpose of CE is that it provides a human-friendly information representation format that is directly processable by machine agents with a clear and unambiguous underlying semantics [9]. In this sense it is the direct equivalent of existing technical languages such as XML, or more specifically OWL/RDF when considering the explicit semantic meaning. Since CE is simply an alternative representation format for ontologies and corresponding data it is directly compatible (through simple translation) with extant and ongoing model development activities such as SensorML [10] and the W3C Semantic Sensor Network Incubator Group [11]. All of the CE examples used in this paper are thus directly processable by machine agents and we believe more consumable by human readers than non-CNL equivalent technical representations. The improvement of CE syntax to allow further linguistic variety and expressivity without undermining the unambiguous semantic grounding is a topic of current research.

CE is used to define the conceptual model that underpins the domain in question. These sentences take the form of concept and relationship definitions (via “conceptualise” sentences) and the definition of logical inference rules (not shown in this paper). Once the conceptual model is defined subsequent assertions can be made according to these concepts and relationships, for example “there is a...”. Extensive working examples are given throughout this paper.

Concepts (defined by `conceptualise` sentences) may be specialisations of other concepts (indicated by `is a` declarations). Relationships may be defined between concepts (for example, the relationship `provides` between the concepts `asset` and `capability`). The following sample definitions are from the asset ontology (see Figure 3):

```
conceptualise an asset type A
  ~ is rated as ~ the NIIRS rating R and
  ~ provides ~ the capability C.

conceptualise a ~ system type ~ S that
  is an asset type.

conceptualise a ~ sensor type ~ S that
  is a system type.
```

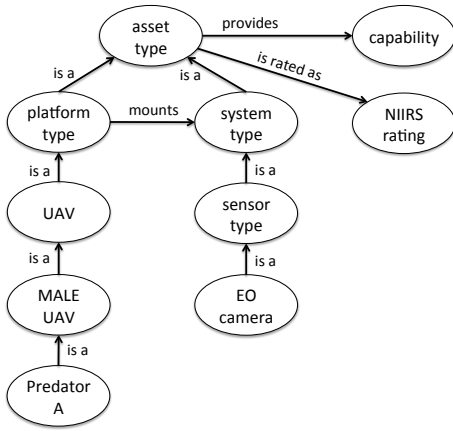


Fig. 3. Graphical depiction of part of the CE asset ontology

conceptualise a "platform type" P that
is an asset type.

conceptualise the platform type P
"mounts" the system type S .

conceptualise a "UAV" U that is a platform type.

conceptualise a "MALE UAV" M that is an UAV.
Note: MALE = Medium Altitude, Long Endurance.

conceptualise a "Predator A" P that is a MALE UAV.

conceptualise an "EO camera" E that is a sensor type.
Note: EO = Electro-optical.

Sentences beginning with 'Note:' are annotations. The types of system that can be mounted on a platform, including sensors, are specified using the `mounts` relationship, for example:

there is a MALE UAV named 'MALE UAV platform type' that
mounts the sensor type 'EO camera sensor type' and
mounts the sensor type 'TV camera sensor type' and
mounts the sensor type 'FLIR camera sensor type' and
mounts the sensor type 'LADAR sensor type'.

Here, a "prototypical" instance MALE UAV platform type is used to capture all of the "prototypical" instances of sensor types that can be mounted on the platform, for example EO camera sensor type.

As described in detail in [12], we automate the assignment of sensing assets to tasks using a knowledge base derived from the NIIRS method of rating imagery data [13]. Thus, NIIRS ratings are associated with assets in the above definitions. Also following the NIIRS approach, we allow intelligence tasks to be defined in terms of basic capabilities such as *detect*, *identify*, and *distinguish*, and one or more kinds of object of interest which we call *detectables*. A task is defined as follows:

conceptualise the task T
"requires" the intelligence capability IC and
"is looking for" the detectable thing DT and
"operates in" the spatial area SA and
"operates during" the time period TP and
"is ranked with" the task priority PR .

The definition includes a spatial area-of-interest, a time period, and a priority to allow tasks to be ranked if assets are scarce). Example task instances will be shown in Section VI.

The NIIRS approach allows automatic matching of tasks

to asset capabilities, by means of encoding NIIRS knowledge in machine-processable form. Here is an example knowledge base intelligence clause in CE, for the case that wheeled vehicles can be identified with visible imagery at NIIRS rating 4 or better:

there is an intelligence clause named `ic003` that
fulfills the intelligence capability `identify` and
is looking for the detectable thing 'wheeled vehicle' and
provides the capability 'visible sensing' and
is rated as 'visible NIIRS rating 4'.

NIIRS ratings are associated with platforms and sensor types, by means of the `provides` relationship:³

there is an EO camera named 'EO camera sensor type' that
provides the capability 'visible sensing'.

there is a Predator A named 'Predator A platform type' that
is rated as the NIIRS rating 'visible NIIRS rating 6' and
is rated as the NIIRS rating 'RADAR NIIRS rating 4'.

From these ontology and instance declarations, we perform automated matching of tasks to available assets (as part of the DCPD direction stage) using the following procedure:

- 1) A task instance t is specified by either a user or the system.
- 2) The intelligence capability and detectable things associated with t are used to select a set of relevant NIIRS clauses, C .
- 3) The combinations of NIIRS capability instances associated with the clauses in C are matched against the potential combinations of platform type and sensor type consistent with the `mounts` relationship.
- 4) The choice of suitable asset is thus constrained by these combinations of platform type and sensor type; it is further constrained by the assets available in the area-of-interest at the time specified for task t .⁴

As an example, a task to detect wheeled vehicle detectables requires the capabilities visible sensing and visible NIIRS rating 4 by the example intelligence clause above. The first of these capabilities is provided by EO camera sensor type and the latter by Predator A platform type. There is a `mounts` relationship between these, so they constitute a deployable platform-sensor combination. The asset catalogue described in Section III will determine if an instance of this platform-sensor combination is available for the task.

A key feature of this approach is that the user states their task in terms of *what* their information requirements are, rather than *how* those requirements should be fulfilled. The NIIRS-based knowledge base determines all potential ways of achieving the requirements. For example, identification of vehicle types can be achieved under particular circumstances by imagery of various kinds (visible, radar, infra-red, multi-

³As with the `mounts` example above, instances of the `provides` relationship are defined between "prototypical" instances of the relevant sensor and platform types.

⁴The final selection of an asset is made using a utility function as described in [8] but left outside the scope of this paper.

spectral) and also acoustic sensing [14].⁵

So far, we have focussed on the selection of assets for tasks. Once an asset is assigned to collect data of a particular kind (for example, location reports for vehicles of a particular kind in a particular region) the collected data will be pre-processed to present an information feed of an appropriate type. Often this will be CNL. For example, vehicles of particular kinds may be localised by imagery or acoustic data. In either case, data processing services can extract the key features from the “raw” data, and make these available as a feed of CNL reports. This is appropriate since the task was expressed in terms of what information is required (vehicle identifications): users expect to receive intelligible reports rather than “raw” imagery or acoustic data — though they may wish to access the original data subsequently. Processing the collected data to yield CNL reports greatly facilitates the fusion of hard and soft data; we provide some examples in the next two sections.

V. USING CNL IN EXPLOITATION OF SOFT SOURCES

In the D2D context, there are large volumes of unstructured data generated from intelligence reports, web commentary and other soft information sources. There is a significant challenge in extracting and properly representing information from such soft sources so that it can be used to support decision-making activities. This needs to happen rapidly when a situation is unfolding at a high tempo. Moreover, soft and hard information typically needs to be combined on-the-fly to provide the best-available intelligence picture of the situation. We see CNL as a key element of a solution here, as it offers an approach to unstructured information processing that facilitates human/machine interaction through a common readable language. The aim is to exploit the synergies of people and machines working together to more efficiently extract task-relevant information at a higher fidelity of representation.

We propose that the use of a CNL in this way facilitates clearer communication between humans when discussing information presented by the system, and also enables the system to act directly on the information without the need to transform to/from another technical representation. The ability for the machine to reason on the CE, and to communicate the rationale for the reasoning, and for the human user to contribute relevant new information in the same format we believe provides a strong unifying representational layer.

The CE based approach that can be used in support of Natural Language Processing is very similar to Ontology Based Information Extraction [15] in that the underlying CE model is used as the basis against which to run the NLP extraction process for the target corpus. A differentiator in our approach is that the underlying CE model is augmented with lexical information to express the ways in which the concepts and relationships in the model are typically expressed in natural language corpora, and this knowledge is used by

the system to attempt to identify salient information from the corpus and output the results in the form of CE sentences.

Our example vignette deals with common kinds of objects in surveillance/policing scenarios: people, vehicles, and places. We have already seen how some of these are captured in the task-related ontology of “detectable” things in the previous section. By using existing NLP techniques, we can extract information on such objects and their features from unstructured text messages. To consider a simple example of a message that might be received from the patrol in step 1 of our vignette:

“Suspicious vehicle driving south: black saloon car with license plate ABC123”

We can automatically extract from this text an observation represented as an instance of the `vehicle` concept in CE:

```
there is a vehicle named v01253 that
  has 'black saloon car' as description and
  has black as colour and
  has saloon as body type and
  has ABC123 as registration.
```

The full extraction would include additional information about location, direction, and the origin of the report (in this case a patrol). The CE syntax has specific extensions to handle common types of meta-data information, including *certainty*. In this vignette this could be applied in a number of contexts, for example, the certainty associated with information generated from natural language processing, or the certainty that a task allocation is accurate.

The CE data would be fed to any analyst who has requested such data as part of their information requirements, and stored for subsequent retrieval and processing. Such a use of NLP constitutes an early part of the processing DCPD stage, and feeds into the remainder of the cycle as follows:

- Further processing can be performed — either automatically or with the intervention of an analyst — on the extracted CE information; for example, fusing it with information collected from other (hard or soft) sources.
- CE messages are already in a relatively convenient form for dissemination, being human-readable.
- Information extracted in CE from soft sources can be used automatically to generate further information requests, automating the transition from dissemination to direction.

Since NLP is inherently imprecise it is important that the NLP processing agent is able to communicate contextual information about the processing result when required. For example this may take the form of an associated certainty related to information arising from NLP (see above), or may be additional information regarding failed parses or other such exceptions that can be used to trigger actions such as requesting a human review, or making a default tasking assumption.

VI. DETAILED WALKTHROUGH OF THE VIGNETTE

Based on the CE elements introduced in the previous two sections, we now provide a detailed walkthrough to illustrate how CE can facilitate hard/soft information fusion and automate much of the DCPD cycle for D2D activities. The

⁵The public domain version of NIIRS from which we have derived our knowledge base is restricted to imagery data; however in principle the knowledge base can be extended where other kinds of data are applicable to achieve a given task.

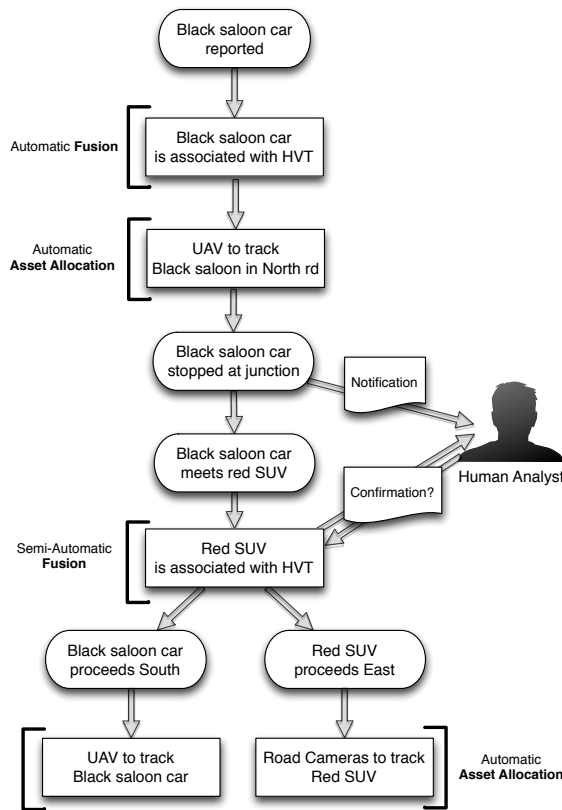


Fig. 4. Walkthrough of the vignette

sequence of activities is summarised in Figure 4, highlighting when fusion is performed automatically by the system and when human intervention is required to infer new data. The output of these fusion steps is then used by the system to perform automatic asset allocation.

Step 1

The patrol on North Road issues the following semi-structured text message: “Suspicious vehicle driving south: black saloon car with license plate ABC123”. A CE processing service uses NLP techniques to extract the following information in CE form:

```

there is a vehicle named v01253 that
  has 'black saloon car' as description and
  has black as colour and
  has saloon as body type and
  has ABC123 as registration.

```

Additional information about location, direction and reporting patrol are also stored, but not shown here. The processing service now uses the features of this message to query stored sources for related information. CE is retrieved which determines that the vehicle with registration ABC123 is known to be associated with a high-value target (HVT) John Smith.

```

the person p670467
  is known as 'John Smith' and
  is a high value target and
  has ABC123 as registered vehicle.

```

New information is automatically stated as a result (shown as *automatic fusion* in Figure 4):

```

there is a HVT sighting named h00453 that
  has the vehicle v01253 as target vehicle and
  has the person p670467 as hvt candidate.

```

An important feature of the CE-based approach is that human-readable *rationale* explaining the inference is also generated (typically using CE inference rules, not shown here):

```

there is a HVT sighting named h00453 because
  the person p670467 is a high value target and
  the person p670467 has ABC123 as
    linked vehicle registration and
  the vehicle v01253 has ABC123 as registration.

```

The production of an HVT sighting instance automatically triggers the generation of an information requirement, in the form of a task instance (an example of an automatic and high tempo processing-dissemination-direction chain):

```

there is a task named t327893 that
  requires the intelligence capability localize and
  is looking for the vehicle v01253 and
  operates in the spatial area 'North Road' and
  is ranked with the task priority high.

```

NIIRS-based reasoning determines that this task can be solved by, amongst other things, a MALE UAV equipped with an EO camera sensor. The UAV in the area of North Road is the only suitable asset available, and it is assigned to the newly-generated task (shown as *automatic asset allocation* in Figure 4). All of this is done without intervention by a human analyst; a CE description of the new task may be posted to the analyst for their information but, depending on their preferences, they would not necessarily be alerted at this point.

Step 2

The UAV locates and starts to track the HVT’s car, posting CE updates to the analyst’s board, for example:

```

there is a tracking report named tr04657 that
  has the vehicle v01253 as target and
  has the person p670467 as candidate hvt and
  has moving as current status and
  is located at the spatio-temporal point loc59695.

```

This report places the black saloon at a specific place and time (spatio-temporal point). After a period of time, the UAV produces a message that the vehicle has stopped:

```

there is a tracking report named tr04658 that
  has the vehicle v01253 as target and
  has the person p670467 as candidate hvt and
  has stopped as current status and
  is located at the spatio-temporal point loc69543.

```

Here, the spatio-temporal point loc69543 corresponds to the known location ‘Central Junction’⁶. At this point we assume that the analyst wishes to be automatically alerted to this significant change (they are interested when an HVT reaches a potential rendezvous point) and calls up the most recent imagery from the UAV, showing that the HVT’s black saloon has stopped next to a red SUV. The analyst tags the image with CE, indicating the red SUV is of interest:

⁶The location “Central junction” has been annotated with a spatial area, and the proximity or containment of the current location of the vehicle (loc69543) is computed by an existing spatial database function.

```
there is a vehicle named v01892 that
  has red as colour and
  has SUV as body type and
  is associated with the vehicle v01253.
```

Note that the location and source analyst details are also saved in CE along with a link to the image, and the association would contain more than just a link between the two vehicles (not shown here). We now have a second vehicle of interest (because of its indirect association with the HVT) so an automatic tasking request is issued similar to the one in Step 1. With the vehicles co-located in the same area, the UAV is able to take this new task in addition to its original task.

Step 3

At this point, we assume that the UAV has not been able to obtain imagery to allow a service to identify the registration of the red SUV. However, the analyst's tagging of the red SUV as being an object of interest allows the system to match this to messages generated earlier from the roadside camera system on Western Road:

```
there is a vehicle named v01879 that
  has 'red SUV' as description and
  has red as colour and
  has SUV as body type and
  has XYZ789 as registration.

there is a vehicle sighting named vs04514 that
  observed the vehicle v01879 and
  has east as heading and
  is located at the spatio-temporal point loc92453.
```

Based on an automated system search for any further information about the red SUV, these messages are shown to the analyst, along with a link to the associated imagery which can be requested if desired. The analyst manually checks for disconfirming information and concludes that the likelihood is that this single report is the red SUV in question. We refer to this as *semi-automatic fusion* in Figure 4 as the system and analyst work cooperatively. The analyst now associates the two sightings with the CE sentence:

```
the vehicle v01879 is the same as the vehicle v01892.
```

This "is the same as" assertion means that all the properties of each instance are propagated across, thereby linking the previously detected registration number to the SUV from Central Junction that the analyst wishes to track.⁷ This sentence is the product of information gathered from the UAV, previously-stored observations from the roadside camera, and the analyst's own intuition (for example, based on a search for additional relevant or disconfirming information in the same timeframe and reading of rationale behind system generated inferences). When making "is the same as" assertions in CE, a human or machine agent making the assertion can also record their certainty or assumptions using the CE syntax.

Step 4

The two vehicles now begin moving, the black saloon proceeding south, and the red SUV heading east. Our two tracking tasks are now:

⁷The assumed registration may subsequently be confirmed when the vehicle is tracked further.

T1: track the black saloon in the region of South Road;

T2: track the red SUV in the region of Eastern Road.

Unfortunately, the UAV cannot be assigned to both the South Road and Eastern Road areas, so it must drop one of the two tasks, T1 or T2. NIIRS-based reasoning determines that the roadside camera system on Eastern Road is suitable for identifying the red SUV when it passes that point. No other asset is available in the South Road area. The allocation system (described in detail in [8]) determines the best assignment (again, shown as *automatic asset allocation* in Figure 4):

- assign the UAV to task T1;
- assign the Eastern Rd camera system to task T2.

Step 5

The asset assignment procedure allows a degree of lookahead. The assigned camera provides only limited utility to the task of tracking the vehicle east. Once it goes beyond the camera's restricted range, further assets will need to be engaged. Prospective tasks can be generated ahead of time as part of a "what if" analysis, to determine if assets would be available to cover the vehicles' potential progress. This is beyond the scope of this paper, but is covered in [16]. This provides further support for the direction stage of the DCPD cycle. If assets are not likely to be available, a further option may be to generate an alert to local law enforcement, to be on the look out for the red SUV.

Discussion

In the walkthrough, we have shown how our approach assists the automation of the DCPD cycle supporting decision-making. We can see multiple iterations of the loop in evidence:

- Step 1: soft information collected, processed, and disseminated from the North Rd patrol causes a direction task to be generated, resulting in the UAV being assigned to collect hard data;
- Step 2: hard data collected and processed from the UAV and disseminated to the analyst prompts them to request more detailed hard data (imagery) from the UAV, hence identifying a new information requirement (to track the red SUV);
- Step 3: the association of the red SUV with the HVT and previously-collected identification information is part of the exploitation aspect of the processing stage, involving fusion of hard and soft information (with analyst input);
- Step 4: the evolving situation requires automatic modifications to the set of information requirements (as part of the direction stage), resulting in a change to the collection assignments;
- Step 5: further iterations of the cycle are likely as the red SUV heads east, including the potential for disseminating an alert to local law enforcement.

Moreover, we have shown how hard/soft fusion is enabled by common ontologies and the uniform representation of information in CE, enabling the creation of new associations:

- the red SUV is associated with the HVT;
- the registration of the red SUV is (probably) XYZ789.

Importantly, in the D2D context, the human analyst remains in the loop, but we avoid overburdening them:

- direction and collection is set up automatically (steps 1, 2, 4, and 5);
- human-readable updates are posted, but the analyst is alerted only when something significant occurs (step 2);
- relevant stored information is automatically retrieved and aligned with new information (steps 1 and 3);
- the D2D system is cooperative: the software makes some associations automatically, the analyst uses their judgement and intuition to make others, and all information is assembled and linked in CE (steps 2, 3, and 5).

A key point to note regarding the CE-based approach is that rationale is available for every step, and can answer questions such as:

- Why is the red SUV of interest?
- Why wasn't the UAV tasked to follow the red SUV?
- Why is the red SUV asserted to have registration XYZ789?

The CE model can be augmented by rules to address the balance of human vs automated processing and decision-making. Examples would include default actions in certain situations, and fall-back steps to be taken if a required human review is not available in a required time window. The specific rules can be tailored to each application and can involve contextual factors such as time of day and level of other tasks underway, as long as these factors are present in the CE conceptual model. Aspects of the system can thus be "tuned" between varying levels of human vs automated processing (although clearly there will always be some decisions that must be taken by a human user in any such system).

VII. CONCLUSION

In this paper, we have shown various roles the use of a controlled natural language can play in improving automation and agility of D2D and DCPD processes, exploiting hard and soft information sources. Communication between the human decision-maker and the system is facilitated by a common understanding of the CNL. In cases where the user needs to work directly with CE, training time can be reduced, compared with formal languages such as XML and RDF/OWL. There are increased options for cooperative working, enabling a flexible mix of automatic and semi-automatic fusion as seen in our vignette. The system can generate traces of its working in CNL that serve as rationale for what it did (or tried to do), and trust between the user and system can be improved by this greater degree of transparency. In our future work, we will examine the generation and processing of more natural sentence structures for CE, ways in which expert users can potentially provide data or metadata to the system and thus extend its capabilities (including the provision of valuable local knowledge at the edges of the network), and conversational modes of interaction between user and system.

ACKNOWLEDGMENT

This research was sponsored by the US Army Research Laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the US Government, the UK Ministry of Defence or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. Geeth de Mel's work was partly carried out while at University of Aberdeen, UK.

REFERENCES

- [1] B. Broome, "Data-to-decisions: a transdisciplinary approach to decision support efforts at ARL," in *Proc Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR III (SPIE Vol 8389)*. SPIE, 2012.
- [2] J. Llinas, C. Bowman, G. Rogova, A. Steinberg, E. Waltz, and F. White, "Revisiting the JDL data fusion model II," in *Proc Seventh International Conference on Information Fusion (FUSION 2004)*, 2004, pp. 1218–1230.
- [3] G. Pearson and T. Pham, "The challenge of sensor information processing and delivery within network and information science research," in *Proc Defense Transformation and Net-Centric Systems 2008 (SPIE Vol 6981)*. SPIE, 2008.
- [4] T. Pham, G. Cirincione, D. Verma, and G. Pearson, "Intelligence, surveillance, and reconnaissance fusion for coalition operations," in *Proc 11th International Conference on Information Fusion*, 2008.
- [5] J. Sowa, "Common logic controlled english," 2004. [Online]. Available: <http://www.jfsowa.com/clce/clce07.htm>
- [6] M. Gomez, A. Preece, M. Johnson, G. de Mel, W. Vasconcelos, C. Gibson, A. Bar-Noy, K. Borowiecki, T. L. Porta, D. Pizzocaro, H. Rowaihy, G. Pearson, and T. Pham, "An ontology-centric approach to sensor-mission assignment," in *Proc 16th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2008)*. Springer, 2008, pp. 347–363.
- [7] J. Wright, C. Gibson, F. Bergamaschi, K. Marcus, R. P. K. Marcus, R. Pressley, G. Verma, and G. Whipples, "A dynamic infrastructure for interconnecting disparate ISR/ISTAR assets (the ITA Sensor Fabric)," in *IEEE/ISIF Fusion 2009 Conference*, June 2009.
- [8] D. Pizzocaro, A. Preece, F. Chen, T. L. Porta, and A. Bar-Noy, "A distributed architecture for heterogeneous multi sensor-task allocation," in *Proc 7th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS'11)*, 2011.
- [9] D. Mott, "Summary of ITA controlled english," 2010. [Online]. Available: <https://www.usukita.org/papers/5658/details.html>
- [10] A. Robin, S. Havens, S. Cox, J. Ricker, R. Lake, and H. Niedzwiedek, "OpenGIS sensor model language (SensorML) implementation specification," Open Geospatial Consortium Inc, Tech. Rep., 2006.
- [11] L. Lefort, C. Henson, and K. Taylor, "Semantic Sensor Network XG final report," 2011. [Online]. Available: <http://www.w3.org/2005/Incubator/ssn/XGR-ssn-20110628/>
- [12] G. de Mel, M. Sensoy, W. Vasconcelos, and A. Preece, "Flexible resource assignment in sensor networks: A hybrid reasoning approach," in *1st International Workshop on the Semantic Sensor Web*, 2009.
- [13] J. M. Irvine, "National Imagery Interpretability Rating Scales (NIIRS)," in *Encyclopedia of Optical Engineering*. Marcel Dekker, Oct. 2003, pp. 1442–1456.
- [14] B. Guo, Y. Wang, P. Smart, N. Shadbolt, M. Nixon, and T. Damarla, "Approaching semantically-mediated acoustic data fusion," in *Proc IEEE MILCOM*, 2007.
- [15] D. C. Wimalasuriya and D. Dou, "Ontology-based information extraction: an introduction and a survey of current approaches," *Journal of Information Science*, vol. 36, no. 3, pp. 306–323, 2010.
- [16] F. Chen, T. L. Porta, D. Pizzocaro, A. Preece, and M. B. Srivastava, "A system architecture for exploiting mission information requirement and resource allocation," in *Proc Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR III (SPIE Vol 8389)*. SPIE, 2012.